



# UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE  
United States Patent and Trademark Office  
Address: COMMISSIONER FOR PATENTS  
P.O. Box 1450  
Alexandria, Virginia 22313-1450  
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/732,003	12/07/2000	Rodney A. DeKoning	99-284	1605

24319 7590 09/08/2003

LSI LOGIC CORPORATION  
1621 BARBER LANE  
MS: D-106 LEGAL  
MILPITAS, CA 95035

EXAMINER

THOMAS, SHANE M

ART UNIT	PAPER NUMBER
----------	--------------

2186

DATE MAILED: 09/08/2003

3

Please find below and/or attached an Office communication concerning this application or proceeding.

# Office Action Summary

Application No.

09/732,003

Applicant(s)

DEKONING ET AL.

Examiner

Shane M Thomas

Art Unit

2186

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

## Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If the period for reply specified above is less than thirty (30) days, a reply within the statutory minimum of thirty (30) days will be considered timely.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133).
- Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

## Status

- 1) ☐ Responsive to communication(s) filed on \_\_\_\_.
- 2a) ☐ This action is FINAL. 2b) ☒ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

## Disposition of Claims

- 4) ☒ Claim(s) 1-15 is/are pending in the application.
- 4a) Of the above claim(s) \_\_\_\_ is/are withdrawn from consideration.
- 5) ☐ Claim(s) \_\_\_\_ is/are allowed.
- 6) ☐ Claim(s) 1-15 is/are rejected.
- 7) ☐ Claim(s) \_\_\_\_ is/are objected to.
- 8) ☐ Claim(s) \_\_\_\_ are subject to restriction and/or election requirement.

## Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☒ The drawing(s) filed on 07 December 2000 is/are: a) ☐ accepted or b) ☒ objected to by the Examiner.
- Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
- 11) ☐ The proposed drawing correction filed on \_\_\_\_ is: a) ☐ approved b) ☐ disapproved by the Examiner.
- If approved, corrected drawings are required in reply to this Office action.
- 12) ☐ The oath or declaration is objected to by the Examiner.

## Priority under 35 U.S.C. §§ 119 and 120

- 13) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some \* c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
2. ☐ Certified copies of the priority documents have been received in Application No. \_\_\_\_.
3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).
- \* See the attached detailed Office action for a list of the certified copies not received.
- 14) ☐ Acknowledgment is made of a claim for domestic priority under 35 U.S.C. § 119(e) (to a provisional application).
- a) ☐ The translation of the foreign language provisional application has been received.
- 15) ☐ Acknowledgment is made of a claim for domestic priority under 35 U.S.C. §§ 120 and/or 121.

## Attachment(s)

- 1) ☒ Notice of References Cited (PTO-892)
- 2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
- 3) ☐ Information Disclosure Statement(s) (PTO-1449) Paper No(s) \_\_\_\_.
- 4) ☐ Interview Summary (PTO-413) Paper No(s). \_\_\_\_.
- 5) ☐ Notice of Informal Patent Application (PTO-152)
- 6) ☐ Other: \_\_\_\_\_.

## **DETAILED ACTION**

### ***Drawings***

The drawings are objected to as failing to comply with 37 CFR 1.84(p)(5) because they do not include the following reference sign(s) mentioned in the description: element 186 of figure 5. A proposed drawing correction or corrected drawings are required in reply to the Office action to avoid abandonment of the application. The objection to the drawings will not be held in abeyance.

### ***Claim Rejections - 35 USC § 112***

The following is a quotation of the second paragraph of 35 U.S.C. 112:

The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the applicant regards as his invention.

Claim 1-7, 14, and 15 are rejected under 35 U.S.C. 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention.

As per claim 1, lines 16-17 are ambiguous and open to multiple interpretations. It is unclear whether Applicant is stating that the "storage spaces" contain both the "logical volume data" and the "selected storage arrays" or if the "volume information" contains information describing the "logical data volume," the "storage spaces" containing the "logical data volume," and the "selected storage arrays." Clarification is required.

Art Unit: 2186

As per claim 4, the phrase "migrating a portion of the one of the storage spaces" is ambiguous and open to multiple interpretations. It is unclear whether the "migration" involves moving the actual data contained in the "storage spaces" of the first storage array to the second storage array or whether the "storage spaces" themselves of the first storage array (where data has yet to be written) are reallocated to the second storage array. The examiner will examine the claim with regard to the latter interpretation: reallocating a portion of the "storage spaces" from the first array to the second *before* the data has been written to the first storage array. It is the duty of the applicant to particularly point out and distinctly claim the subject matter which applicant regards as the invention.

As per claim 14, is it unclear whether the "plurality of storage arrays" are included in the "consolidated storage array," the "host device," or the "storage area network." The examiner will examine the claim with regard to the "plurality of storage arrays" as being included in the "consolidated storage array." Clarification is required.

As per claims 14 and 15, the term "data transfer performance" is ambiguous and open to multiple interpretations. The "data transfer performance" might refer to the bandwidth of each storage array or the total amount of data that has been written to the storage array (physical space). The examiner will examine the claim regarding the "data transfer performance" to be the latter interpretation in that the amount of data loaded (transferred) into the storage array (physical space used) is what is being monitored. It is the duty of the applicant to particularly point out and distinctly claim the subject matter which applicant regards as the invention.

***Claim Rejections - 35 USC § 102***

The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless -

(e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

Claims 1, 2, 5-7, and 13 are rejected under 35 U.S.C. 102(e) as being anticipated by Nguyen et al. (U.S. Patent Application Publication No. US 2002/0004883).

As per claim 1, Nguyen shows in figure 5 a “consolidated storage array” which comprises a plurality of “storage arrays” (TD1-TD5), Connection Block CB0, and a network storage manager (NSM), contained in device controller ACSC. The examiner is referring to this section of the apparatus shown in figure 5 as a “consolidated storage array” because it has the ability to “consolidate” the storage arrays TD1-TD5 into a single logical volume as will be shown. The examiner is referring to TD1-TD5 (tape drives used in the examples of Nguyen) as “storage arrays” since their tape media contain a collection of sequential (“arrayed”) storage locations. A host device, which the examiner is regarding as a Connection Block CB1-CB4 connected to its respective data processor DP1-DP4, runs software to manage the access protocol for the plurality of storage arrays TD1-TD5 (paragraph 9) and utilize a created logical volume. The host devices are connected to the to the “consolidated storage array” via network CN and the fiber network connection N.

As can be seen in the example of paragraph 27, the NSM (network service manager) receives needed storage performance requirements from one of the host devices. The examiner

Art Unit: 2186

is also considering the size of the virtual data volume to be included in the “performance requirements” because the NSM would have to know how much physical space to allocate. Upon receiving the requirements and analyzing the storage arrays’ data rate (performance capacities of TD1-TD4) individually, the NSM constructs a virtual device (or “logical data volume”) by selecting some of the storage arrays so that the combination of arrays meets the required data rate (performance capacity). Further in paragraph 27, the NSM then instructs the respective Connection Block of the host device to stripe the data to storage arrays TD1-TD5 as needed achieving the required data rate. The virtualization or creation of the virtual device – herein referred to as a virtual data volume – is performed by functions in the Connection Blocks (CB1-CB9). These functions can be realized in *software* or hardware (see paragraph 28).

Referring to paragraph 29, if a host device requests access to a virtual data volume, the host sends a request to the RM software of the NSM, which in turn communicates with the ACSLS (maintaining the volume database) to determine if the volume is already created. If new, the RM software makes a request to the RA software to reserve the necessary number of storage arrays, and then instructs the RC software to configure the connection blocks (CB5-CB9) for data transfer from whichever is the requesting data processor. The RC software also instructs the requesting processor’s CB connection that the data will be striped to the allocated number of storage arrays. Once the RM software is informed that everything is set up, it sends a notification to the requesting processor that the virtual device (“logical data volume”) is ready for data transfer. Because the RA software decides which physical devices to utilize in creating the virtual device and to what locations of the physical device(s) the data will be sent (and striped if necessary), it is inherent that the information (“volume information”) regarding the

Art Unit: 2186

physical location of the virtual data volume be contained as an entry in a database apparatus.

(Nguyen states in paragraph 35 that when the RM software receives the request from the host, that it ascertains the physical location of the physical medium(s) that contains the selected data volume – that is if the host wishes to read data already contained in the storage array.) Further, it is inherent that if striping data across the storage arrays is required, that a “striping definition” accompany the volume information – either directly or indirectly (pointer) – since the physical location of a virtual data volume can be ascertained from the “consolidated storage array” (paragraph 35). Finally (returning to paragraph 29), the “consolidated storage array” then sends this volume information to the respective Connection Block of the requesting host device to configure the striping software to write the data in the virtual data volume.

As per claim 2, paragraph 29 (as previously discussed) cites an example of the host device (via data processor DP1) issuing a volume request specifying the data transfer rate, data access time, data compression, etc. to the RM software of the “consolidated storage array.” Volume information regarding which storage arrays and locations are sent to the host’s Connection Blocks CB5-CB9. The data access functions (that can be performed in software – paragraph 28) of the Connection Blocks of the host are then configured to stripe the data from the processor DP1 into the allocated storage arrays. According to the example, CB1 is configured by the RC software to stripe the data to the three different storage arrays TD1-TD3.

As per claim 5, Nguyen states in paragraph 14 and 37, that the “consolidated storage array” dynamically configures the storage arrays responding to the requests of the host system, specifically the data processors. These requests can come as manual operator commands, pre-programmed algorithms, rules, application program initiated requests, and the like. Further, in

Art Unit: 2186

paragraph 15, Nguyen states the rules that can be implemented comprise response time constraints, data file transfer size, file transfer rates, and data file size bounds.

As per claim 6, Nguyen states in paragraph 28 that the [data access] functions in the Connection Blocks of the host can be implemented in software. These functions achieve the virtualization process, or creation of the “logical data volume” by means of striping the data from the host into the allocated storage arrays. The connection blocks are responsible for the striping of the data to the corresponding storage arrays as cited in paragraph 29.

As per claim 7, as has been discussed, data is striped from the processors into a virtual data volume of the storage arrays. Further, the examiner is stating that a “striping definition,” either directly or indirectly, is inherently included with the “volume information” of the virtual data volume that is stored in the ACSLS or another database. Nguyen’s storage system must have a way of accessing the data that has been striped to the storage arrays as well as to determine which location contains the parity information for the stripe.

As per claim 13, the examiner is referring to the entire apparatus of figure 5 of Nguyen as a “storage area network.” The same rejection for claim 1 is applied to lines 1-6 since the rejection discusses obtaining the performance requirements for the virtual data volume (“logical data volume”), and that if the performance requirement exceeds the capacity of an individual array (in the example cited by Nguyen, data rate was the performance requirement), multiple storage arrays could be combined and the data striped to obtain the required data rate. The host device “establishes the parameters” of the virtual data volume in its request to the “consolidated storage array.”

The same rejection for claim 2 is applied to lines 7-11 of claim 13.



Art Unit: 2186

The same rejection for claims 1 and 7 are applied to lines 12-14 of claim 13. Claim 1's rejection discusses creating an entry in a database apparatus that contains "volume information" regarding where the physical location of the data volume is located on the physical medium(s) and that the "striping definition" is inherently included in the "volume information" for accessing the data that has been striped to the storage arrays as well as to determine which location contains the parity information for the stripe.

As per lines 15-17 of claim 13, as was previously discussed in claim 1's rejection, the striping software in the host (contained in functions in the Connection Blocks) uses the "volume information." Once the RC software sets up the data system shown in figure 5 to transfer data, it instructs the Connection Block, corresponding to the requesting processor, to stripe the data to the allocated virtual data volume contained on storage arrays that the RA software has allocated.

Claims 1, 2, 5-8, and 11-13 are rejected under 35 U.S.C. 102(e) as being anticipated by Ito et al. (U.S. Patent No. 6,408,359).

As per claims 1 and 8, the examiner is referring to a "consolidated storage array" as the system of figure 1 combined with the storage arrays composed of storage devices 1807a-1807c and 1807d-1807f, respectively, of figure 5. Virtual data volumes (or "logical data volumes") are distributed equally to each storage array and also to each storage device within the array (column 3, lines 40-53). The examiner is referring to a "host device" as the combination of (1) an external device, which can run software applications for a Video-On-Demand service (column 1, lines 21-24), coupled to the (2) Input/Output Control Unit 1808 (figure 5) which is then coupled to the (3) storage device control units, 1806a and 1806b. The external device sends a "file" (or

Art Unit: 2186

virtual data volume) request that includes information specifying the number (size) of storage locations of the virtual data volume to the Destination Determination Part (1803 of figure 1) of the consolidated storage array (column 3, lines 43-49). The examiner will refer to this information as the “performance requirement” of the virtual data volume. Ito states that the virtual data volume is spread equally throughout the storage arrays and storage devices so that the consolidated storage array can respond to a read request at a high speed, since access requests do not concentrate on specific storage devices nor specific storage array. Therefore, the performance capability (bandwidth and/or physical volume size) of a single array is less than that of the collective storage arrays comprising the consolidated storage array. Further, it can then be seen that the virtual data volumes’ performance requirements exceed the performance capabilities of a single storage array because they span across the storage arrays in the system of Ito.

Ito’s system indirectly analyzes the storage arrays for their performance capabilities (physical volume size, or remaining volume size in this instance) by keeping track of empty storage locations in the Empty Area Managing Part (1802 in figure 1). This part manages the empty “blocks” (storage locations not containing data associated with a virtual data volume) of each of the storage devices using an address of each empty storage location as well as first and second identifiers, which identify a specific storage device and specific storage array, respectively (refer to column 3, lines 19-39).

The Destination Determination Part (1803 of figure 1) selects and configures locations from each of the storage devices of each storage array by utilizing the information provided by the Empty Area Managing Part and the Configuration Managing Part (1801 of figure 1), which

Art Unit: 2186

manages the number of storage devices and respective identifying information of the system.

The flow chart of figure 6 shows a method used to select and configure empty storage locations by selecting a set number of locations per loop cycle (steps S2305-S2307), allocating them for the virtual data volume (step S2308), comparing the number of locations allocated to the number requested in the performance requirement for the volume sent by the host (step S2309), and selecting and configuring more locations if the number allocated is less than the number requested. Thus the method of figure 6 shows how the performance requirements (size of volume requested by the host) are met by the consolidated storage area system in figure 5 (see column 3, lines 19-39).

Volume information is created and managed by the system in the address position file such like element 2401 of figure 7. Each virtual data volume has an address position file associated with it (column 11, lines 11-17). It is inherent that once the virtual data volume is created among the storage arrays and the storage devices that the applications running on a host device, which sent the volume request to the system of Ito, would use the allocated storage locations to store and retrieve data. Address position file (volume information) is sent from the consolidated storage array to the input/output control until (1808 of figure 5) of the host device. The I/O control unit distributes the data inputted from the external device of the host to the storage device control units (1806a and 1806b). These control units then write the data to the allocated storage locations (refer to column 8, lines 54-61).

Regarding claim 8, the examiner now refers to the storage array, composed of storage devices 1807a-1807c and 1807d-1807f, respectively, as a "consolidated storage array" (CSA). Further, the examiner now refers to the system shown in figure 1 of Ito as the "CSA primary

Art Unit: 2186

device.” Since the Configuration Managing Part 1801 of figure 1 stores the number of storage devices and respective identifying information, it is inherent that this Part is coupled the storage arrays – either directly or indirectly – such that the “CSA primary device,” system of figure 1, is connected to the storage arrays at least by the aforementioned connection. Finally, the examiner refers to the collection of the host, consolidated storage array, and CSA primary device as a “storage area network.”

As per claim 2, as has been shown in the rejection for claims 1 and 8 above of Ito, an external device (of a host) issues a file (or virtual volume) create command to the Destination Determining Part of the “consolidated storage array.” Once the volume has been allocated, the CSA sends the address position file (volume information) to the host device to enable the applications running on the host to utilize the allocated volume to store and retrieve data.

As per claims 5 and 11, as has been shown in the rejection for claims 1 and 8 above of Ito, the external system specifies the size requirement for the virtual data volume by the number of storage blocks (locations) it needs in order to sufficiently store its data (see column 5, lines 43-49). Further, as has been stated, the examiner is referring to the “size” of the volume to be the “performance requirement” of the volume.

Regarding claim 11, Ito is referring to a user as something that is utilizing the storage device management system to access (store/retrieve) data. In one example, the system of Ito is being used in a Video-On-Demand system as a video server. In this case, it could be seen that the “user” could be a “client” of the server and is requesting a video object and would therefore read data from Ito’s system. More indirectly, the examiner in interpreting a user to be the person controlling the “client” wishing to access the “videos” contained on the system of Ito.

Art Unit: 2186

As per claim 6, as has been detailed in the rejection for claims 1 and 8 above, the Destination Determining Part of the “consolidated storage array” allocates storage locations within each storage device of each storage array for each request for a virtual data volume. The software application running on the external device of the host sends data through the external device and supplies it to the I/O control unit (1808 of figure 5) of the host which then sends the data, using the address position file (volume information) received from the consolidated storage array, to the respective storage device control units (1806a and 1806b). The examiner is referring to this procedure as “striping” data from the host device since the data is not written a specific storage device nor specific storage array

As per claim 7 and 12, the same definition of striping as defined in the rejection for claim 6 is herein used for claim 7 and 12’s rejection. Further, a “striping definition” is incorporated into the address position file (volume information) that is created and sent from the Destination Determining Part (1803 of figure 1) of the consolidated storage array. This “striping definition” supplies the host’s I/O control unit 1808 with the address of the allocated storage locations of the virtual storage array so it can begin sending the data received from the host’s external device to the respective storage control units 1806, which write the data to the storage devices 1807.

As per claim 13, the same rejection for claim 1 is applied to lines 1-6; the same rejection for claim 2 is applied to lines 7-11. Further, the same rejection of claims 6 and 7 are applied to lines 12-17.

***Claim Rejections - 35 USC § 103***

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

Claims 1, 2, 5-7, and 13 are rejected under 35 U.S.C. 103(a) as being unpatentable in view of Nguyen (U.S. Patent Application Publication No. US 2002/0004883).

As per claims 1 (lines 16-20), 7, and 13, if the applicant can show that the describing of a volume information in a database apparatus is not anticipated by Nguyen, then it would have been obvious to one of ordinary skill in the art at the time the invention was made to create and store volume information, when allocating a virtual data volume, in a database in the NSM, so that after the data had been written (striped if necessary) by the Connection Blocks, it could have been read out correctly when an access to the virtual data volume was made. Further it would have been obvious to include a striping definition – either directly or indirectly (pointer) – with the volume information in order to allow the processor that had requested the data from the virtual data volume to have had known where the parity information for the data would have been located so that the data could have been verified before being read out.

The rejections for claims 2, 5, and 6 follow the same rejections as discussed above in the 35 USC 102(e) rejections under Nguyen, respectively.

Art Unit: 2186

Claims 3 and 4 are rejected under 35 U.S.C. 103(a) as being unpatentable over Nguyen (U.S. Patent Application Publication No. US 2002/0004883) in view of Allen et al. (U.S. Patent No. 5,151,990).

Nguyen does not utilize a method that monitors and reallocates “storage spaces” of a part of a virtual data volume of a storage array. Allen describes in column 4, lines 25-48, a method for comparing the remaining capacity of a volume to a predetermined threshold and signaling an error message if the remaining capacity falls below the threshold. The examiner is regarding the size of the storage array to be the “maximum performance capability” of the array. Further Allen describes utilizing the error message to initiate a responsive reallocation of available space within the system. This feature allows for operations that would normally have to abort to continue running between the time when the error was generated and the end of the operation.

Therefore, it would have been obvious to one having ordinary skill in the art at the time the invention was made to incorporate the monitoring and reallocation of resources system of Allen onto the Network Service Manager (NSM) of Nguyen in order to have prevented an overflow situation where the data being striped from the host to the storage arrays would not have been lost due to a shortage of volume capacity. Referring to claim 4, the examiner is regarding the “first and second” storage arrays to be any two of storage arrays (TD1-TD5) that the RM software has allocated for a request by a host (assuming the host requires a performance rate greater than an individual storage array TD1-TD5).

Claims 8, 11, and 12 are rejected under 35 U.S.C. 103(a) as being unpatentable over Nguyen (U.S. Patent Application Publication No. US 2002/0004883) in view of Applicant’s admitted prior art.

As per claim 8, Nguyen shows a plurality of “storage arrays” (TD1-TD5) in figure 5. However Nguyen does not particularly mention if those arrays each contain a plurality of storage devices. The applicant’s admitted prior art states on page 2, lines 4-9, that storage arrays can consist of multiple storage devices, which can be accessed in parallel for a much greater overall bandwidth. Therefore, multiple locations of a volume of data, contained within the storage array on different storage devices, can be accessed simultaneously. Further the applicant states on page 2, lines 12-18, that more than one of these “storage arrays” can be combined together in a storage array system and can be accessed in parallel, generating a much greater overall bandwidth and transaction rate than a single storage array. A data volume can then be divided up and allocated to more than one of the storage arrays of the storage array system in order to achieve the desired bandwidth and transaction rates for access to the data volume.

Therefore, it would have been obvious to one having ordinary skill in the art at the time the invention was made to combine multiple storage devices into the storage arrays of Nguyen (the storage arrays being previously discussed in claim 1’s rejection) in order to gain the added bandwidth, transaction rate, and storage capacity of the larger storage array system. Nguyen could have then read and written data to the “consolidated storage array” (being previously discussed in claim 1’s rejection as well) at a much higher transaction rate, thus speeding up the overall bandwidth of the data system of figure 5. In addition, the host system of Nguyen would have had a larger virtual data area, thereby allowing the system to have utilized more virtual data volumes. Nguyen states in paragraph 41 that his claimed invention could have been realized with removable storage devices or fixed media devices.



Art Unit: 2186

The rejection for claim 8 follows Nguyen's claim 1 rejection. The examiner is regarding the "storage area network" to be the entire apparatus of Nguyen as shown in figure 5. The examiner is regarding the ACSC, containing the Network Storage Manager, combined together with TD1 to be the consolidated storage array's "primary device" since the NSM, as has been detailed in claim 1's rejection, is connected to the storage arrays to the Fiber Channel Network N of figure 5 via CB0; executes volume create software based on requests from a host device that analyze the storage array data rates individually; constructs (if necessary) a virtual device to meet the data rate performance requirement of the requesting host (data processor and its respective Connection Block) (paragraph 27); and configures the virtual data volume to accept and distribute data from the host throughout the storage arrays.

As per claim 11, Nguyen states in paragraph 14 that the RC software of the NSM configures the data storage resources [during the creating of a virtual data volume] and responds to manual operator (user) commands, pre-programmed algorithms, rules, application program initiated requests and more. Further, in paragraph 15, Nguyen states that the rules can comprise response time constraints, data file transfer size, data transfer rates, data transfer bounds, and more.

As per claim 12 with regard to Nguyen, the same rejection as claim 7 applies. Additionally, it could have been seen that the striping software of Nguyen, contained in the function of the Connection Blocks (refer to paragraph 28), could have been exchanged with the striping and redundancy techniques admitted as prior art by the applicant (page 2, lines 18-19). Since the striping of Nguyen would have taken place within the requesting host's Connection Block (refer to paragraph 29), the requesting processor would not need to realize the physical to

Art Unit: 2186

virtual translation. Thus, striping data from the host across the storage devices in different storage arrays could have been realized.

Claims 9 and 14 are rejected under 35 U.S.C. 103(a) as being unpatentable over Nguyen (U.S. Patent Application Publication No. US 2002/0004883) in view of Applicant's admitted prior art in further view of Burgess et al (U.S. Patent No. 5,796,633).

As per claim 9, since the storage arrays of modified Nguyen are arrays of multiple storage devices, the examiner is regarding an entire storage array as a logical volume – since it spans multiple physical storage devices. Further, the examiner is regarding that the “performance capability of the logical volume be either the size (amount of data that can be stored on the array) or the frequency of access (reads/writes) to the storage device, and the “maximum performance capacity” to be the maximum amount of data that can be stored in the array or the maximum frequency of access the array can withstand. Modified Nguyen does not utilize a method that monitors the performance of the storage arrays to determine whether the arrays are performing within a predetermined range of the maximum performance capacity. Burgess shows performance monitoring relating to different performance parameters or thresholds of physical and logical volumes such as frequency of reads and writes (column 8, lines 66-67). These predetermined thresholds can be determined by the user (column, 7, lines 1-3). The user can instruct a configuration file to monitor current memory loads, available physical memory, available page files, and the available virtual memory size (column 6, line 4-9). Further the monitoring system of Burgess can alert the user or a log file if a monitored performance counter dips below its predetermined threshold or “maximum performance capability.” An alert thread can monitor the percentage of free space remaining on each logical

Art Unit: 2186

volume (column 7, lines 9-16). The monitoring and warning system of Burgess would determine if a logical volume (storage array) was reaching its maximum storage capacity or the frequency of accesses (read/write) operations. The alert signal could prevent the system from saving further data to the array and be used to trigger a response or action from the system operator to overcome the problem (column 2, lines 45-49). Burgess further states that the functions of the monitoring method could be realized in software (column 15, lines 3-6).

Therefore, it would have been obvious to one of ordinary skill in the art at the time the invention was made to further alter the data storage system of Nguyen to utilize the monitoring method of Burgess by incorporating the method as a further software piece in the NSM ("CSA primary device"). This alteration would have allowed for alerts to be flagged when the logical volume capacities of the storage arrays exceeded a threshold determined by the system operator. These alerts would have allowed an operator to give immediate attention to the performance problem. The monitoring of logical volume capacity would have helped to warn of impending overflow situations where the data being striped from the host to the storage arrays could be lost due to a shortage of volume capacity. Burgess' monitoring method is further beneficial to the system of Nguyen because the performance monitoring can be handled automatically without human intervention (column 2, lines 35-39).

As per claim 14, the same rejection for claim 9 applies. The examiner is regarding the "data transfer performance" of the storage arrays being monitored as being the frequency of the accesses (reads/writes).

Claims 10 and 15 are rejected under 35 U.S.C. 103(a) as being unpatentable over Nguyen (U.S. Patent Application Publication No. US 2002/0004883) in view of Applicant's admitted

Art Unit: 2186

prior art in further view of Burgess et al. (U.S. Patent No. 5,796,633) in further view of Komachiya et al. (U.S. Patent No. 6,571,314).

Modified Nguyen does not disclose a means or method to allow a system operator or the data system itself (figure 5 of Nguyen) to compensate if an alert had been raised signifying that a storage array is within its threshold and maximum capacity or capability. Komachiya teaches optimization of a storage system, where the frequency of accesses to a storage area is over a predetermined threshold, in order to improve the capacity efficiency of the system by migrating the data contained in a single storage area over multiple storage areas. Further, Komachiya teaches that the system can be re-optimized when the frequency of access drops below a predetermined threshold to combine data across multiple storage areas into a single storage area (refer to column 4, lines 60-64 and column 6, lines 43-55). Therefore, it would have been obvious to one having ordinary skill in the art at the time the invention was made to combine the modified system of Nguyen with the teaching of Komachiya in order to have been able to migrate a virtual data volume contained in a single storage array, that was being frequently accessed, to multiple storage arrays, thereby increasing the "capacity efficiency" (or data transfer performance requirement) of the virtual data volume and the overall system. Hence it could have been seen, that the monitoring system of modified Nguyen, coupled with the teaching of Komachiya, would have been able to migrate data across multiple storage arrays in order to have overcome the performance degradation of the system when the frequency of accesses (reads/writes) would have exceeded a predetermined threshold, as would have been defined by the system operator. Here the examiner is referring to the storage area performing near its performance threshold as the "first storage array" and any storage array that data migrates to as a

Art Unit: 2186

result to be the “second storage array.” Also, the examiner is regarding the “size” (number of storage locations) of the virtual data volume to be a “performance capability” of the volume. The modified system of Nguyen would have taken into account the size of the virtual data volume in light of the available size of a “second” storage array when migrating the volume since the system would obviously not have migrated data to a storage array with an inadequate amount of available storage locations.

As stated in the rejection for claim 14 with regard to Nguyen, the examiner is regarding the “data transfer performance” being monitored as the amount of data loaded (transferred) into the storage array (physical space used) or the frequency of the accesses (reads/writes).

The addition of the teaching of Komachiya would have been obvious because it would have allowed the system to migrate data from a storage array that was close to its maximum performance threshold immediately once the data transfer (access frequency) was detected over the predetermined threshold of the storage array. Thus, coupled with the automatic alert monitoring system of Burgess and data migration teaching of Komachiya, the maintenance burden of the system operator would be significantly reduced.

Claims 3, 4, 9, 10, 14, and 15 are rejected under 35 U.S.C. 103(a) as being unpatentable over Ito (U.S. Patent No. 6,408,359) in view of D’Errico et al. (U.S. Patent No. 6,314,503).

As per claims 3, 4, 9, 14, and 15, Ito states that although data is distributed to all storage devices in each “storage array” of the embodiment described in the above rejections, an alternate embodiment could have data distributed to specific storage devices in each “array” (column 11, lines 65-67). Ito does not disclose in either embodiment a method for monitoring the data

Art Unit: 2186

transfer performance, which the examiner is referring to as a “performance capability,” of the storage arrays in order to determine whether the storage arrays are performing within a predetermined range of the maximum “performance capability” of each array. Additionally, Ito does not disclose a method to alleviate the condition of an array performing within its predetermined maximum capability (or maximum data transfer) by migrating a portion of the data between the over-performing (“first”) storage array and another (“second”) storage array.

D’Errico teaches in column 3, lines 49-63, that in a system with multiple storage devices a performance condition can be detected (hence monitored) and alleviated by the re-distribution of system data between the multiple storage devices. Specifically, D’Errico teaches that in a system with a plurality of storage devices, a method comprises the steps of: (A) detecting a segment in the storage system that is accessed frequently and sequentially (a virtual volume) and is stored on one of the plurality of storage devices and (B) in response to step (A), splitting the large data segment into at least two smaller data segments that can be accessed in parallel from at least two of the plurality of storage devices, thereby improving the performance of the overall system.

Therefore, it would have been obvious to one of ordinary skill in the art at the time the invention was made to incorporate the teaching of D’Errico to the data storage system of Ito, in order to alleviate the burden of an array that is performing within its predetermined maximum performance capability by reallocating the virtual data volume that are causing the increase in the data transfer performance of the array, by migrating a portion of the virtual data volume to at least a second storage array. Specifically D’Errico’s method should have been incorporated into the consolidated storage array’s “primary device” (system of figure 1 of Ito) since it performs all

Art Unit: 2186

of the managing procedures of the storage arrays and virtual data volumes, as has been described above. The teaching of D'Errico would have allowed the system of Ito to migrate frequently accessed data (data with a detected high transfer performance) and have distributed it among other storage arrays and devices; thus, an increase in the amount of data that could have been accessed by an application using the data storage system of Ito in the same amount of time would have been achieved. Further, the splitting of the virtual data volume among the storage arrays and devices would be transparent to the host (and therefore the application running on the host). This aspect is advantageous since the method of D'Errico could have been performed automatically without requiring modification to the application running on the host device, and without requiring manual intervention [by the a system operator] (refer to column 6, lines 60-67 thru column 7, line 15). Therefore, it could have been seen that the teaching of D'Errico would have reduced the burden of the data system by increasing the rate at which data could have been accessed, thus improving system performance by splitting a virtual data volume among storage arrays and storage drives while the division would have remained transparent to the host that would have utilized the data system of Ito.

Regarding claim 10 with respect to Ito in view of D'Errico, the examiner is referring to the maximum data transfer performance as the "maximum capacity" of a storage array. Further, the examiner is considering the performance capability of the storage arrays to be the size of the storage arrays and the performance requirement of the "logical data volume" (virtual data volume) to be the size (number of storage locations) of the volume. Therefore, it would have been obvious to have seen that when the predetermined range of the data transfer performance had been crossed for a "first" storage array, the teaching of D'Errico would have compared the

Art Unit: 2186

sizes of the remaining storage arrays in light of the size of the virtual data volume in order to find a "second" (or more) storage array(s) with enough empty (unallocated) storage locations to accommodate the portion of the virtual data partition that had been split.

### ***Conclusion***


Prior art made of record and not relied upon is considered pertinent to applicant's disclosure in PTO-892. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Shane M Thomas whose telephone number is (703) 605-0725. The examiner can normally be reached on M-F 8:30 - 5:30.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Matt M Kim can be reached on (703) 305-3821. The fax phone numbers for the organization where this application or proceeding is assigned are (703) 764-7239 for regular communications and (703) 764-7239 for After Final communications.

Any inquiry of a general nature or relating to the status of this application or proceeding should be directed to the receptionist whose telephone number is (703) 305-3900.



Shane M Thomas  
September 2, 2003



DAVID L ROBERTSON  
PRIMARY EXAMINER